

AD-A188 744

DTIC ACCESSION NUMBER

LEVEL

PHOTOGRAPH THIS SHEET

INVENTORY

AFWA1-TR-87-1178

DOCUMENT IDENTIFICATION

DEC 1987

This document has been approved
for public release and sale; its
distribution is unlimited.

DISTRIBUTION STATEMENT

ACCESSION FOR

NTIS GRA&I ☒

DTIC TAB ☐

UNANNOUNCED ☐

JUSTIFICATION

BY

DISTRIBUTION /

AVAILABILITY CODES

DIST

AVAIL AND/OR SPECIAL

A-1

DISTRIBUTION STAMP

QUALITY
INSPECTED
2

DTIC
ELECTE
FEB 09 1988
S D
CE

DATE ACCESSIONED

DATE RETURNED

88 2 05 098

DATE RECEIVED IN DTIC

REGISTERED OR CERTIFIED NO.

PHOTOGRAPH THIS SHEET AND RETURN TO DTIC-DDAC


AD-A188 744

AFWAL-TR-87-1178

THE ROLE OF ARTIFICIAL INTELLIGENCE IN THE INTEGRATION
OF REMOTELY SENSED DATA WITH GEOGRAPHIC INFORMATION
SYSTEMS

David M. McKeown, Jr.

Carnegie-Mellon University
Computer Science Department
Pittsburgh, PA 15213-3890



December 1987

Interim

Approved for Public Release; Distribution is Unlimited

AVIONICS LABORATORY
AIR FORCE WRIGHT AERONAUTICAL LABORATORIES
AIR FORCE SYSTEMS COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO 45433-6543

NOTICE

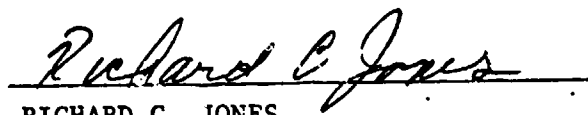
When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility or any obligation whatsoever. The fact that the Government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication, or otherwise in any manner construed, as licensing the holder, or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed by the Office of Public Affairs (ASD/PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

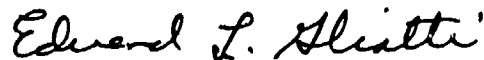


CHAHIRA M. HOPPER
Project Engineer



RICHARD C. JONES
Ch, Advanced Systems Research Gp
Information Processing Technology Br

FOR THE COMMANDER



EDWARD L. GLIATTI
Ch, Information Processing Technology Br
Systems Avionics Div

If your address has changed, if you wish to be removed from our mailing list, or if the addressee is no longer employed by your organization please notify AFWAL/AAAT, Wright-Patterson AFB, OH 45433-6543 to help us maintain a current mailing list.

Copies of this report should not be returned unless return is required by security considerations, contractual obligations, or notice on a specific document.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS			
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.			
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE						
4. PERFORMING ORGANIZATION REPORT NUMBER(S) CMU-CS-86-174			5. MONITORING ORGANIZATION REPORT NUMBER(S) AFWAL-TR-87-1178			
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University	6b. OFFICE SYMBOL (if applicable)		7a. NAME OF MONITORING ORGANIZATION Air Force Wright Aeronautical Laboratories AFWAL/AAAT-3			
6c. ADDRESS (City, State, and ZIP Code) Computer Science Dept Pittsburgh PA 15213-3890			7b. ADDRESS (City, State, and ZIP Code) Wright-Patterson AFB OH 45433-6543			
8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (if applicable)		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F33615-84-K-1520			
8c. ADDRESS (City, State, and ZIP Code)			10. SOURCE OF FUNDING NUMBERS			
			PROGRAM ELEMENT NO. 61101E	PROJECT NO. 4976	TASK NO. 00	WORK UNIT ACCESSION NO. 01
11. TITLE (Include Security Classification) The Role of Artificial Intelligence in the Integration of Remotely Sensed Data with Geographic Information Systems						
12. PERSONAL AUTHOR(S) David M. McKeown, Jr.						
13a. TYPE OF REPORT Interim	13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) 1987 December		15. PAGE COUNT 42	
16. SUPPLEMENTARY NOTATION						
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)			
FIELD	GROUP	SUB-GROUP				
19. ABSTRACT (Continue on reverse if necessary and identify by block number)						
<p>Although databases for geographic information systems (GIS) have been developed to manage digital map data, the integration of remotely sensed imagery and other collateral non-map information is rarely performed. For the most part, the use of sophisticated intelligent spatial databases, in which the user can query interactively about map, terrain, or associated imagery, is unknown in the GIS and cartographic community. In standard GIS systems, the ability to formulate complex queries requiring dynamic computation of factual and geometric properties is severely limited, often reflecting its origin as collections of thematic map overlays. Spatial database research requires the integration of ideas and techniques from many disciplines such as computer graphics, computational geometry, database methodology, image analysis, photogrammetry, and artificial intelligence. In</p>						
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified			
22a. NAME OF RESPONSIBLE INDIVIDUAL Chahira M. Hopper			22b. TELEPHONE (Include Area Code) (513) 255-7865		22c. OFFICE SYMBOL AFWAL/AAAT-3	

UNCLASSIFIED

Block 19 (Continued)

this paper we discuss some ideas on how the scope of geographic information systems can be expanded by utilizing techniques from the AI community which may remedy deficiencies in user interfaces, spatial data representation and its utilization. We draw on ongoing research at Carnegie Mellon University for examples of these techniques in the areas of image/map database and knowledge-based image interpretation.

Table of Contents

Abstract	1
1. Introduction	1
1.1. Integrating Imagery	2
1.2. AI in GIS	4
1.3. Terrain Classification Systems	5
2. Expanding the Scope Of Geographic Information Systems	6
2.1. Query Primitives	6
2.2. Constraining Search	7
3. An Overview of MAPS	9
3.1. A Schema-Based Representation For Spatial Entities	10
3.1.1. Concepts With Multiple Roles	14
3.2. Hierarchical Decomposition	14
3.2.1. Hierarchical Containment Tree	16
3.2.2. Hierarchical Search	18
3.3. Unifying Spatial and Factual Queries	19
3.4. Integrating Multiple Databases	21
4. Application of MAPS In Computer Vision and Graphics	23
4.1. A Geodetic Frame Of Reference	23
4.2. Rule-Based Systems In Image Analysis	24
4.3. Map-Guided Image Segmentation	25
4.4. Airport Scene Interpretation	27
4.5. Stereo Image Analysis	27
4.6. Computer Image Generation	29
5. Conclusions	35
6. Acknowledgements	35
7. References	35

List of Figures

Figure 3-1:	MAPS: System Overview	9
Figure 3-2:	MAPS: Concept and Role Schemata Definitions	11
Figure 3-3:	Concepts from 'washdc' CONCEPTMAP Database [partial list]	12
Figure 3-4:	Conceptmap Database Dictionary: System and User Defined Attributes	13
Figure 3-5:	Conceptmap Database Dictionary: Subrole Attributes	13
Figure 3-6:	Multiple Roles for 'george washington university'	15
Figure 3-7:	CONCEPT for 'kennedy center'	16
Figure 3-8:	MAPS: Hierarchical Spatial Containment	17
Figure 3-9:	MAPS: Containment Tree Entry for Theodore Roosevelt Island	19
Figure 3-10:	Query Operations on Role Schemata	20
Figure 3-11:	Template Match Query Attributes	20
Figure 3-12:	Multiple CONCEPTMAP Database Organization	21
Figure 3-13:	Search In Multiple Databases	22
Figure 4-1:	MACHINESEG: Map-Guided Feature Extraction	26
Figure 4-2:	SPAM: System Organisation	28
Figure 4-3:	STEREOSYS: Stereo Verification	30
Figure 4-4:	WASH3D: Computer Image Generation	31
Figure 4-5:	WASH3D: Terrain Scene of Washington D.C. With CONCEPTMAP Database Overlay	32
Figure 4-6:	WASH3D: Scene of Washington D.C. With CONCEPTMAP Cultural Features	33
Figure 4-7:	WASH3D: Closeup of Foggy Bottom Area With CONCEPTMAP Terrain and Cultural Features	34

Abstract

Although databases for geographic information systems (GIS) have been developed to manage digital map data, the integration of remotely sensed imagery and other collateral non-map information is rarely performed. For the most part, the use of sophisticated intelligent spatial databases, in which the user can query interactively about map, terrain, or associated imagery, is unknown in the GIS and cartographic community. In standard GIS systems, the ability to formulate complex queries requiring dynamic computation of factual and geometric properties is severely limited, often reflecting its origin as collections of thematic map overlays. Spatial database research requires the integration of ideas and techniques from many disciplines such as computer graphics, computational geometry, database methodology, image analysis, photogrammetry, and artificial intelligence. In this paper we discuss some ideas on how the scope of geographic information systems can be expanded by utilizing techniques from the AI community which may remedy deficiencies in user interfaces, spatial data representation and its utilization. We draw on ongoing research at Carnegie Mellon University for examples of these techniques in the areas of image/map database and knowledge-based image interpretation.

1. Introduction

In this paper we discuss some possible roles for Artificial Intelligence (AI) in the design of future geographic information systems. The remote sensing and geographic information communities have had significant interactions in the use and analysis of spatial data, whether sensed or derived, since their inception. Other research disciplines that have contributed in this area include photogrammetry, geodesy, earth sciences, cartography, and computer science. Because of this diversity the term 'geographic information systems' has a variety of meanings, often shaded by the background or application area of the researcher. This is easily seen in the GIS literature; see, for example, the proceedings of the ASPRS AUTOCARTO Workshop series, or¹, a conference on Geographic Information Systems in Government sponsored by the U.S. Army Engineer Topographic Laboratories. This paper is written from the perspective of a computer scientist working the areas of computer vision and artificial intelligence as applied to spatial databases and knowledge-based aerial image analysis. The view of a spatial database is quite broad and general: it should contain access to three-dimensional entities, both natural and manmade, tied to their geographic coordinates, with a rich and flexible set of attributes and descriptions. A flexible attribute set is one where the user, rather than the system architecture, defines the sets of properties or attributes that a spatial entity can assume. This also holds for the retrieval mechanisms provided by the system. This broad scope has some advantages. We are not focused on a particular application domain such as forestry, agronomy, geology, or cartography. Our goal is to produce a general framework within which any application can be accommodated. Further, these requirements begin to cross the line between traditional database methodologies and knowledge base design and maintenance methodologies². How well we achieve that goal can be judged by the variety of

applications that can be supported by the spatial database architecture, some of which are briefly described in this paper.

The primary thrust of this paper is that the concept of a geographic information system should be extended to include remotely sensed imagery, and that such imagery can become a focus for user interaction, spatial database update, and emerging image interpretation and analysis capabilities. Secondly, sub-areas of AI such as knowledge representation and utilization, knowledge acquisition, and computer vision have reached a stage of maturity that will have impact on spatial database representation and land use classification systems based on remotely sensed imagery. In this Section we expand upon these ideas.

1.1. Integrating Imagery

There are several reasons why remotely sensed imagery should be an integral component of a geographic information system and why the lack of such an image display component has caused an apparent dichotomy between the producers of spatial data and traditional GIS users. That is, GIS users do not have the tools to be producers of spatial knowledge, and they are not generally explicitly aware of the source or accuracy of their data. First, most GIS databases are produced by digitizing map products. The original map products may be produced from imagery, using rigorous photogrammetric techniques, or may be derived from other cartographic products. There may be several levels of abstraction and generalization between the cartographic product and the remotely sensed data that was originally used to produce it. Once a map product is produced, none of the associated data used to produce it remains available. That is, the product is divorced from the quality, reliability, and timeliness of the source material. For example, interpolation of digital terrain models from contour maps is not unheard of, nor is interpolation of discrete depth points for bathymetric models. How accurate is this interpolated data? Only by knowing the source and derivation methods can we begin to establish data accuracy.

Second, the user interfaces for most geographic information systems are limited to:

- Display of all topographic data within an area of interest. Usually this is accomplished by scaling and drawing all polygon vectors within a bounding rectangle.
- Recall of topographic data based upon the data type. For example, all polygons labeled as 'forest' in the terrain coverage overlay. Often the user is provided a small fixed length field within which he is free to store 'attributes' without any structuring or consistency checking by the GIS.
- Search for internally generated identifiers.
- Display of graphic icons within an area of interest.

While these techniques may suffice in a variety of applications we believe that there are significant advantages in using remotely sensed imagery as a medium of interaction in a geographic information system. The user can formulate queries into the spatial database using multi-resolution imagery to specify an area of interest for spatial indexing, or to specify a generic type of feature by pointing to an exemplar in the scene. With the additional capability of explicit image-to-map correspondence by dynamically calculating ground coordinates as user indicates image (or display) coordinates, the image can act as a map for a user unfamiliar with an area of interest.

This should not be surprising. It is often the case that the image contains more detail and information to a user than a digitized map. It may be easier for a user to orient himself to a geographic area using imagery than a map overlay. Further, given the basic capability of image-to-map and map-to-image correspondence there can be interesting juxtapositions: very small scale imagery can be used to show overall spatial organization at a level of political, or geophysical boundaries which then switch to parcel level descriptions on large scale imagery. Pointing to a low resolution overview image in one portion of the display can cause high resolution imagery of the corresponding area to be dynamically displayed. Finally, given the appropriate photogrammetric and graphic analysis tools, the user can modify the baseline database to suit his own particular needs. One can then view the GIS as a general topologically consistent framework within which the user can:

- Make minor modifications by observing changes or missing features in the update imagery.
- Intensify local areas of a large global database without having to maintain a homogeneous level-of-detail.
- Add features of temporal interest to a baseline GIS and 'playback' changes over time.

There are several reasons why this has not been done within the GIS or computer vision communities.

- Cost: requires a high resolution, multi bit per pixel display to be effective.
- Requires an photogrammetric model associated with each image so that geographically referenced data can be accurately superimposed on the image and that coordinates generated via user interactions with the display can be transformed from display space, to image space, to geographic coordinates.
- Requires a partial solution to the 'clutter problem', the ability to dynamically prune the superimposition of objects on the image in areas where there is high object density.
- Requires dynamic search of a potentially large database to decide what features to display based (primarily) on their spatial location and intrinsic properties.

In Section 3 we describe the MAPS system developed at Carnegie Mellon in which we have experimented with many of these issues, both in the context of user interfaces to spatial data, and as a component of image analysis and interpretation systems.

1.2. AI in GIS

AI can play a role in the development of advanced geographic information systems. If we view geographic information systems as an application area for the use of AI technology and approaches, there are several sub-areas of AI which may hold promise such as knowledge representation and utilization, knowledge acquisition, and computer vision. Some task areas where AI techniques may find utility are:

1. Non-traditional representations of factual information based on semantic networks of information. Meta-knowledge about the relationships of objects, normally reserved for application-specific systems should become more closely linked with the topological properties and attributes of spatial objects and their relation to other spatial objects.
2. Categorizing and unifying knowledge about multi-spectral classification of specific types of remotely sensed imagery, such as LANDSAT MSS and Thematic Mapper, and SPOT. Such knowledge is usually tightly coupled with systems which perform image analysis. A more general methodology would be to separate this knowledge from the data interpretation systems which use it. The spatial database then becomes a common repository for general knowledge which is accessible from a variety of applications.
3. Integration of image analysis systems and other data sources with spatial databases so that the spatial database becomes a source of knowledge utilized during image analysis or classification as well as the repository of the partial and final results of the analysis.

The expectations that high performance knowledge intensive systems based upon AI technology will transform the remote sensing analysis community by solving many of the still difficult problems in automated landuse classification including improving their accuracy and robustness are perhaps premature. Current knowledge-based systems typically exhibit a rather narrow character -- often described as *shallow*. Though substantial knowledge is collected in the rules, capable of being released appropriately to perform the task, there is no ability to reason further with that knowledge. The basic semantics of the task domain are not understood by the system. New rules are not learned from experience nor does behavior with the existing rules become automatically tuned. To point out such limitations is not to be hypercritical of the current art. Indeed, the important scientific discovery behind the success of AI knowledge-based expert systems is precisely that sufficient bodies of such shallow knowledge could be assembled, without any of the supporting reasoning and understanding ability, and still prove adequate to perform

real consultation tasks in the medical and industrial worlds. There are many reviews of 'expert systems' in the literature, among those^{3, 4} describe systems that have actually been implemented as well as some analysis of their capabilities and level of development. Some interesting comments on the state-of-the-art and future directions can be found in McDermott⁵.

1.3. Terrain Classification Systems

One of the most attractive areas for integration of AI methods with remotely sensed imagery lies in analysis and interpretation. Recently, researchers investigating and building terrain classification systems^{*} have adopted an "AI" approach. These traditional classification systems have been recast to use "rules" in place of decision trees or multispectral feature analysis. This is unfortunate because the lack of depth of knowledge, the primary cause of poor performance in the original formulation, has not been addressed. It is naive to believe that by merely changing methodologies, from decision trees to "rules" or "evidential reasoning", that task performance will improve due to the different representation or classification methodology.

What is lacking in many terrain classification systems is a depth of knowledge applied to the image under analysis, analogous to the "shallowness" of reasoning in most first-generation knowledge-based systems. For some tasks simple mappings of spectral properties to landuse categories may suffice. However, as our performance expectations rise and the complexity and detail of the remotely sensed image increase, it is unlikely that classification systems based on shallow reasoning can be pushed much beyond current performance levels. In fairness, these techniques generally are driven by, and may be limited by, the inherent low resolution of the remotely sensed imagery from sources such as Landsat MSU, Thematic Mapper, and SPOT. Nevertheless, world and domain knowledge must be applied to the classification problem in a systematic method, rather than using *ad hoc* scene-by-scene adjustments if we expect to improve the state-of-the-art in terrain classification.

A brief description of the differences between image processing and image understanding may shed some light on how this can be accomplished. Traditional image processing preforms a transformation of a source image to produce another image. Examples are 'difference images' for change detection, 'thresholded images' for foreground-background segmentation, and 'labeled images' for pixel-based classification. Image understanding attempts to produce a symbolic representation of the contents of the scene in terms of non-image descriptions. Examples are 3-dimensional stereo reconstruction, various 'shape

* Included in terrain classification are traditional landuse and landcover classification systems developed for mapping and earth resources such as forestry and agronomy.

from' techniques, model-based interpretation of aerial imagery, and autonomous navigation.

Image processing generally does not bring much domain knowledge to bear on the transformation process, while image understanding attempts to use scene geometry, illumination cues, domain constraints, and spatial constraints to construct a symbolic description of the structures in the scene, usually using multiple (non-image) levels of representation. Is there an analogy between image processing and landuse classification? We believe that there is a need for the development of more knowledge intensive symbolic processing in landuse classification. One role for AI may be in the development of 'landuse understanding' systems which utilize knowledge about climate, soil composition, terrain, and general *a priori* map knowledge, in addition to the spectral information contained in remotely sensed imagery. Such systems will require a spatial database in which such knowledge can be stored and retrieved when appropriate. We believe that current GIS organizations are too limited for such a task. In the following section we outline some capabilities that would move us closer to this goal.

2. Expanding the Scope Of Geographic Information Systems

One viewpoint for comparing the capabilities of geographic information systems is to focus on the ability to perform various types of spatial and factual queries. The emphasis on query capability can be used to drive other functional capabilities related to data representation, retrieval, and display. It is also interesting to note that in our experience there is little difference in the underlying functionality between queries generated by an interactive user, and those generated by an application program using the spatial database. Those differences which do arise tend to be ones that are tied to economy of large transactions in the application versus small grain responses to an interactive user. A second viewpoint for comparison is how search is performed as a consequence of a spatial or factual query. For large GIS it is not reasonable to assume that brute-force search for factual or spatial relations will provide tolerable performance. Nor can we assume that all relationships can be known (precomputed) at the time a query is formulated. This implies a set of space-time tradeoffs between directly representing spatial relationships and their computation upon demand. In this section we discuss several types of queries that should form a basis for an advanced GIS and some methods for constraining search during queries.

2.1. Query Primitives

For the purposes of discussion, let us define a spatial database system, or *image/map database*, as one which integrates digital map, image, and spatially indexed factual data. A *spatial query* on this heterogeneous database can be performed by some combination of the following query primitives:

1. Pointing to the digital display of a remotely sensed image or map using a cursor, mouse, or other pointing device.
2. Providing geographic coordinates and spatial constraints such as proximity and containment.
3. Enumerating the properties of various man-made or natural objects.
4. Naming the spatial entity.
5. Indicating a 'level of detail' for the query.
6. Invoking an image classification or analysis system to find a specified feature or class of objects.

Primitives (1) and (2) can be characterized as *signal access* into the image/map database. That is, a transformation into or search within the geographic component of the database must be performed. Primitives (3), (4), and (5) can be characterized as *symbolic access* into the image/map database. Search within the symbolic descriptions of each spatial entity must be performed. Primitive (6) is a *hybrid access*, requiring signal access to the image under analysis, and symbolic access to properties of the features or objects being extracted.

Primitive (3) has been performed by static representation of a small number of spatial relationships such as 'intersection' and 'containment', usually within a relational database model. However, relationships such as 'within radius' and 'adjacency' are not amenable to such precomputation since their semantics vary widely depending on the task at hand. Further, traditional relational models do not have adequate methods for structuring large amounts of spatial data to avoid (nearly) complete search of the database. We believe that an image/map database should provide computation on demand of spatial relationships with frequently accessed facts stored after this computation in a 'spatial cache' or 'memo file'. The overall goal is to structure the database and the computational geometry primitives so that a search is performed to enumerate those features which are potentially effected by the query, followed by their efficient computation. This partitioning is critical to all forms of knowledge access, whether signal or symbolic, but is particularly important in the computation of spatial relationships.

2.2. Constraining Search

In order to realize acceptable performance for dynamic queries within a large spatial databases the search associated with each query must be bounded by the system. Search can be limited in several ways. *Depth limiting search* implies that there is a hierarchical relationship between spatial entities, and that all spatial entities have an intrinsic scale below which they become invisible to the search process. *Breadth limiting search* implies that large sections of the database can be ignored by observing that the root of a subtree of the hierarchical database is disjoint (not contained within nor intersects) from the query area. Breadth limiting search is commonly implemented in a variety of regular spatial

decomposition schemes such as quadtrees and k-d trees, and non-regular methods such as hierarchical containment trees. However, depth limiting search is an interesting new idea for constraining search in spatial databases. It also provides a practical implementation for the representation of 'level of detail' required for primitive (5).

One method to achieve a depth limiting search in the image/map database is to explicitly represent the notion of a 'level of detail' associated with each query. The 'level of detail' can be provided by the user as a portion of his query specification or can be defaulted to the limits imposed by the underlying data. For example, given that the user is pointing at a 1:1,000,000 scale digitized map, the system can enforce that point queries will be limited to searching the database for features which would be portrayed on a 1:50,000 scale map. Thus, individual buildings, roads, and small-scale features that might be represented in the image/map database would not participate in the search. The user should be able to describe the level of detail using terms like 'to the county level', 'to the neighborhood level', and 'to the finest detail'. These terms can be converted by the system to scaling factors.

One extension to giving the user explicit control over the level of detail of his query is to add the capability of "continuous spatial zooming". By this we mean that the system should have the ability to interpolate or extrapolate spatial knowledge in order to provide a response to the spatial query should information not be stored at the requested level of detail. How this might be performed for certain types of spatial data composed of discrete samples such as elevation grids or terrain coverage databases as well as for imagery is fairly well understood. However, for symbolic data this requirement has problems similar to those found in cartographic generalization. This implies that the spatial database maintains a model of what information is directly known, what information can be derived, and the limits of its derivation capability. We believe that this is an interesting area for future research which requires interaction between AI and GIS.

In the remainder of this paper we expand on these themes and give some concrete examples from our experiences in building research systems for image/map databases, knowledge-based image interpretation, and three-dimensional (3D) scene generation. In Section 3 we briefly describe the MAPS system, an image/map database with many of the characteristics and goals outlined in previous sections. In Section 4 we discuss several past and current research systems developed at Carnegie Mellon, with emphasis on how each one utilized spatial knowledge to perform its task. Those components of the MAPS system that were particularly relevant to the overall system design are described. The reader is directed to the bibliography for a more detailed discussion of each of these systems.

3. An Overview of MAPS

The MAPS spatial database^{6,7,8} was developed between 1960-1984 supported by the DARPA Image Understanding Program as research into large-scale spatial databases and spatial knowledge representation. It is interesting that this system has expanded from its original research goal of developing an interactive database for answering spatial queries into a component of several knowledge-based image understanding systems under development at Carnegie Mellon University. MAPS is a large-scale image/map database system for the Washington D.C. area that contains approximately 200 high resolution aerial images, a digital terrain database, and a variety of map databases from the Defense Mapping Agency (DMA). MAPS has been used as a component for an automated road finder/follower, a stereo verification module, and a knowledge-based system for interpreting airport scenes in aerial imager. In addition, MAPS has an interactive user query component that allows users to perform spatial queries using high resolution display of aerial imagery as a method for indexing into the spatial database. This capability to relate, over time, imagery at a variety of spatial resolutions to a spatial database forms a basis for a large variety of interpretation and analysis tasks such as change detection, model-based interpretation, and report generation.

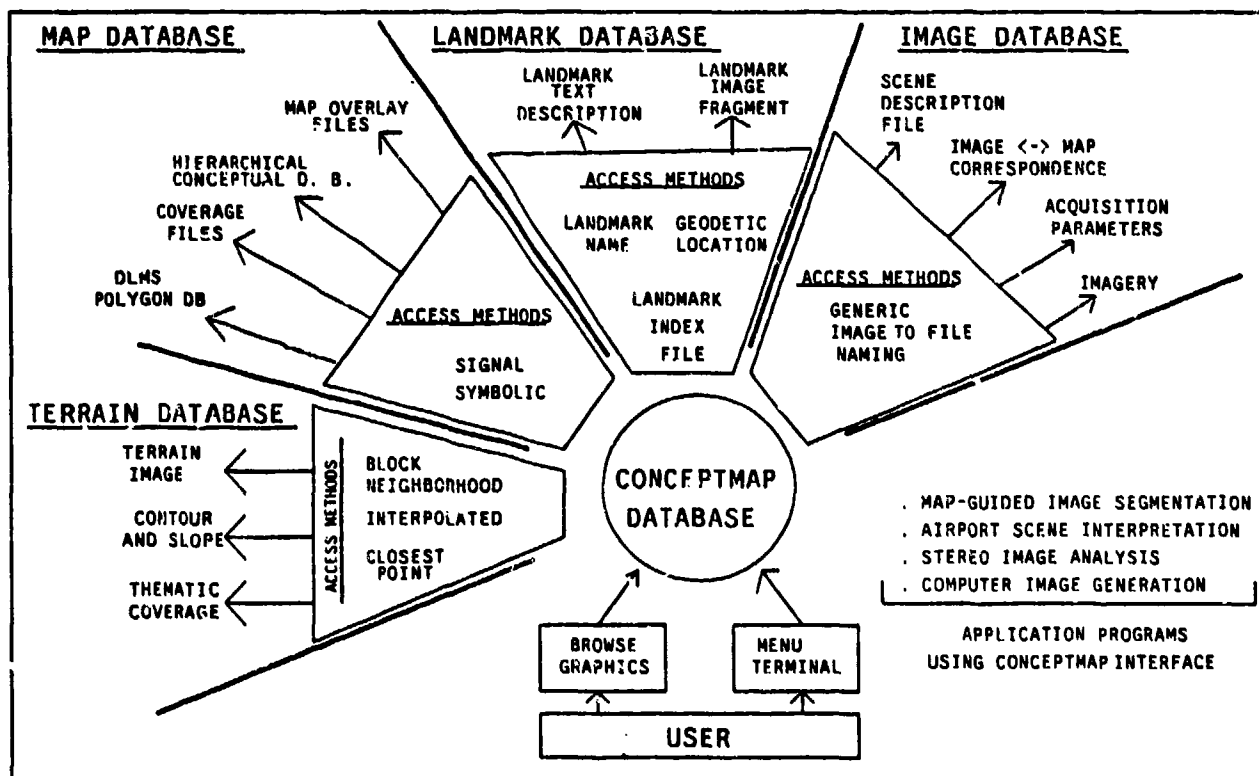


Figure 3-1: MAPS: System Overview

Figure 3-1 shows the system organization of MAPS. Four databases are maintained within MAPS: a digital terrain database, a map database, a landmark database, and an image database. A fifth database, CONCEPTMAP, consists of a schema-based representation for spatial entities and a set of procedural methods that provide a uniform interface to each of the four component databases for interactive users or application programs. It is this interface that allows us to represent and access image, map, terrain, and collateral data in a manner that best suits the intrinsic structure of the data. At the same time the CONCEPTMAP database provides uniform access to a variety of spatial data independent of the particular internal structure. This is in sharp contrast to methods proposed for uniform representation of image and cultural data such as raster data sets and regular decompositions such as quadrees or k-d trees. In the following sections we touch on some interesting aspects of the CONCEPTMAP database.

3.1. A Schema-Based Representation For Spatial Entities

The CONCEPTMAP database uses a *schema-based* representation for spatial entities. Using schemas (or frames) is a well understood AI methodology for representing knowledge. Such a representation can be combined within several problem-solving methods such as semantic networks, scripts or production systems to construct a problem-solving system⁹. Each entity in the CONCEPTMAP database is represented by one *concept* schema and at least one *role* schema. A *concept* can represent any spatial object and associates a name with a set of attributes stored in the *concept* and *role* schemata. Figure 3-2 gives definitions of the slot names for *concept* and *role* schemata. Figure 3-3 gives an partial list of the concepts in the MAPS WASHDC database.

There are three unique identifiers generated by the CONCEPTMAP system which allow for indirect access to additional factual properties of concept or role schemata.

- The *concept-id* is unique across all concepts in all CONCEPTMAP databases. That is, given a concept-id one can uniquely determine the name of the spatial entity.
- The *role-id* uniquely determines a role schema across all CONCEPTMAP databases.
- The *role-geographics-id* uniquely determines a collection of points, lines or polygons in vector notation. Each point is represented as $\langle \text{latitude}, \text{longitude}, \text{elevation} \rangle$.

As shown in Figure 3-2 these identifiers are also used to index into other components of the MAPS database. For example, the *concept-id* is used to search for landmark descriptions of measured ground control points used during the calculation of transform functions for image-to-map and map-to-image correspondence. The *role-id* is used as the basic entity when building a hierarchy tree decomposition as described in Section 3.2. The

GENERAL SCHEMA DEFINITION		ROLE SCHEMA DEFINITION	
SLOT VALUE LIST OF SLOT VALUES SYSTEM GENERATED IDENTIFIER		ROLE-ID ROLE-NAME ROLE-SUBNAME ROLE-CLASS ROLE-TYPE ROLE-DERIVATION ROLE-MARK LIST OF USER-DEFINED-SLOTS LIST OF VALUES FOR USER-DEFINED-SLOTS ROLE-GEOGRAPHICS-ID	
CONCEPT SCHEMA DEFINITION			
CONCEPT-NAME CONCEPT-ID PRINCIPAL ROLE LIST OF ROLE-IDS LIST OF ROLE-PRINTNAMES			
SYSTEM GENERATED IDENTIFIERS		INDEX INTO SPECIALIZED DATABASES	
CONCEPT-ID	ROLE PRINT-NAMES LANDMARK		
ROLE-ID	PROPERTY LIST GEOMETRIC QUERY LIST TEXT HISTORY HIERARCHICAL DECOMPOSITION		
ROLE-GEOGRAPHICS-ID	SPATIAL RELATIONSHIPS (MEMO FILES) 2D SHAPE DESCRIPTION 3D DESCRIPTION CONVEX HULL IMAGE SEGMENTATION IMAGE COVERAGE		

Figure 3-2: MAPS: Concept and Role Schemata Definitions

role-geographics-id is used to acquire the unique geographic position for a *role schema* as well as for linkage into the MAPS image database and segmentation files generated by human interaction or machine segmentation. There are three reasons for this approach. First, it allows CONCEPTMAP to handle very large databases with a minimal amount of information resident in the application process. The identifiers provide a level of indirection to the actual data, which is stored in a variety of formats and may or may not be present for a large subset of the database. Second, we can achieve a great deal of flexibility and modularity in processes which communicate about spatial entities. Given the name of a CONCEPTMAP database, a *concept-id* or *role-id* uniquely determines the entity in question. This facilitates the construction of application programs with simple query structures, requiring a minimum of communication overhead. Finally, given this decoupling from the CONCEPTMAP database, each of the MAPS component databases, image database, terrain database, landmark database, and map database may be physically resident on a different workstation or mainframe.

CONCEPT1	tidal basin	CONCEPT196	1 enfant plaza
CONCEPT2	district of columbia	CONCEPT196	forrestal building
CONCEPT3	northwest washington	CONCEPT197	east potomac park
CONCEPT4	mcmillan reservoir	CONCEPT198	folger library
CONCEPT5	southwest washington	CONCEPT199	senate office building
CONCEPT6	northeast washington	CONCEPT200	visitors center
CONCEPT7	virginia	CONCEPT201	capital hill park
CONCEPT8	maryland	CONCEPT202	capitol plaza park
CONCEPT9	kennedy center	CONCEPT203	mall ice rink
CONCEPT10	ellipse	CONCEPT204	federal office building 6
CONCEPT11	washington circ's	CONCEPT205	natural history museum
CONCEPT12	state department	CONCEPT206	federal aviation administration
CONCEPT13	executive office building	CONCEPT207	freer gallery
CONCEPT14	white house	CONCEPT208	smithsonian institution
CONCEPT15	treasury building	CONCEPT209	george mason memorial bridge
CONCEPT16	department of commerce	CONCEPT210	group hospital building
CONCEPT17	arlington memorial bridge	CONCEPT211	fisher auditorium
CONCEPT18	rfk stadium	CONCEPT212	doctors hospital
CONCEPT19	museum of history and technology	CONCEPT213	route 1
CONCEPT20	key bridge	CONCEPT214	dulles airport
CONCEPT21	kutz bridge	CONCEPT215	rock creek park
CONCEPT22	george mason bridge	CONCEPT216	constitution pond
CONCEPT23	fort stanton reservoir	CONCEPT217	georgetown reservoir

Figure 3-3: Concepts from 'washdc' CONCEPTMAP Database [partial list]

There are three levels of attribution available to users within CONCEPTMAP:

- *system-wide* attributes: stored in role schema.
- *user-defined* attributes: stored in role schema.
- *property-list* attributes: stored in property list database.

CONCEPTMAP allows users to define additional attributes, called *user-defined*, similar in function to the *role-name* and *role-subname* slots described above. Finally, *property-list* attributes can also be defined by the user and are capable of representing a variety of datatypes including 'strings', 'integers', 'double', and 'list' using a simple data structure based on lists of the following:

<'attribute-name' , 'attribute-value'>

Attributes of all three classes are interpreted by CONCEPTMAP using a database dictionary defined for each class type. CONCEPTMAP can be easily configured for a particular application such as geology or forestry simply by developing an appropriate database dictionary. *User-defined* and *property-list* attributes can be defined dynamically by a user at an interactive session.

Figure 3-4 gives a partial dictionary of the *system-wide* slots and representative values for a CONCEPTMAP database. Figure 3-5 is a partial dictionary of *role-subname* values associated with *role-name* values in Figure 3-4. An example of *property-list* attributes is given in Figure 3-11. In the following sections we illustrate some of the power of the *concept/role* schemata for representation and query.

ROLES: UNKNOWN BUILDING BRIDGE ROAD RESERVOIR AIRPORT RESIDENTIAL AREA INDUSTRIAL AREA UNIVERSITY PARKS SPORTS COMPLEX	ROLE-CLASS: UNKNOWN INDUSTRIAL RESIDENTIAL TRANSPORTATION NATURAL FEATURE GOVERNMENT CULTURAL FEATURE COMMERCIAL RECREATIONAL EDUCATIONAL	ROLE-DERIVATION: UNKNOWN HAND-SEGMENTATION MACHINE-SEGMENTATION TERMINAL-INTERACTION LANDMARK-DESCRIPTION DLMS-EXTERNAL UNKN-EXTERNAL
ROLE-TYPES: UNKNOWN PHYSICAL CONCEPTUAL AGGREGATE-PHYSICAL AGGREGATE-CONCEPTUAL	ROLE-MARK: UNKNOWN NONE GEO-QUERY TEMPLATE-QUERY NEW-CONCEPT NEWROLE MODIFY-CONCEPT MODIFY-ROLE NEW-3D MODIFY-3D MODIFY-NAME-ROLE EXTRACT-FROM-DATABASE	USER DEFINED SLOTS: USER-DEFINED 'COMPOSITION' (SLOT 0) UNKNOWN STONE/BRICK SOIL COMPOSITION ASPHALT EARTHEN WORKS CONCRETE ROCK METAL

Figure 3-4: Conceptmap Database Dictionary:
System and User Defined Attributes

ROLE: BUILDING UNKNOWN OFFICE BUILDING GOVERNMENT BUILDING DORMITORY CONCERT HALL MUSEUM PERFORMING ARTS COMPLEX RAILROAD STATION ADMINISTRATION MEMORIAL MEDICAL CENTER BOATHOUSE APARTMENTS HOTEL/MOTEL LIGHT MANUFACTURING GYMNASIUM HOSPITAL LIBRARY CLASSROOMS STUDENT UNION	ROLE: AIRPORT UNKNOWN COMMERCIAL MILITARY	ROLE: UNIVERSITY UNKNOWN DORMITORIES ATHLETIC FACILITIES UNIVERSITY CAMPUS RESEARCH FACILITY ADMINISTRATION STUDENT CENTER CAFETERIA ADMISSION
ROLE: ROAD UNKNOWN INTERSTATE HIGHWAY STREET AVENUE RURAL ROAD TRAFFIC CIRCLE INTERSECTION INTRACITY HIGHWAY ACCESS ROAD	ROLE: BRIDGE UNKNOWN RAILROAD PEDESTRIAN AUTOMOBILE	ROLE: SPORTS COMPLEX UNKNOWN STADIUM BOAT MARINA ICE SKATING RINK OPEN AREA GOLF COURSE
ROLE: PARKS UNKNOWN PLAYING FIELD OPEN AREA POND FORESTED AREA ZOO FORMAL GARDEN BLEACHERS SCULPTURE GARDEN	ROLE: POLITICAL UNKNOWN STATE COUNTY CITY DISTRICT	ROLE: RESIDENTIAL AREA UNKNOWN SINGLE FAMILY HOUSING APARTMENT COMPLEX MIXED HOUSING

Figure 3-5: Conceptmap Database Dictionary:
Subrole Attributes

3.1.1. Concepts With Multiple Roles

Each concept in the CONCEPTMAP database has one concept schema and at least one role schema associated with it. Multiple roles for a concept gives CONCEPTMAP the ability to represent flexibly spatial knowledge. Multiple roles can be used to denote different *views* of the same concept rather than creating multiple independent concepts. If there are multiple roles for a concept, then the *principal role* slot contains the name of the default or preferred view. Each role can define the same spatial area by sharing *role-geographic-id* slots. For example, the concept 'washington d.c.' has three roles with *roles-name* of 'political entity', 'geographic entity', and 'demographic entity' that share the same *role-geographic-id*.

Multiple roles are also used to indicate whole-part relationships. Spatial entities composed of collections of individual objects, where each object is not a separately named concept, can be grouped under an overall concept. Figure 3-6 illustrates this for a subset of the buildings on the campus of 'george washington university'. There is one concept for the university, and a variety of roles including dormitories, gymnasiums, classrooms and student unions. The principal role has a *role-geographic-id* which points to a $\langle \text{latitude/longitude/elevation} \rangle$ description of the perimeter of the campus. The *concept-id* can be used to retrieve *role print-names* for uniquely named buildings, represented as roles, such as 'lisner hall' and 'gwu medical center'. From a practical standpoint this technique avoids a proliferation of *concept-name* values of the form 'gwu hospital', 'gwu classroom 1', 'gwu classroom 2', etc. Since the specification of a *concept-name* is a predominant method for query in CONCEPTMAP, multiple roles group together related spatial entities while allowing for full expression of their individual traits within the database.

Figure 3-7 is a display generated by an interactive user requesting a dump of the concept 'kennedy center'. It illustrates the use of multiple roles to define two different *views* of a concept, and how linkage of one view to the landmark database within MAPS can be performed. The first role schema defines 'kennedy center' as a 'concert hall' with a full three-dimensional description and shape properties derived from its two-dimensional boundary. The second role defines 'kennedy center' to have a *role-derivation* of 'landmark-description'. When this role is accessed within CONCEPTMAP, a linkage to the landmark database is automatically performed using the *concept-id* slot and the textual information associated with the landmark database is displayed.

3.2. Hierarchical Decomposition

The CONCEPTMAP database can be used to build a *hierarchy tree* data structure which represents the whole-part relationships and spatial containment of map feature descriptions. This tree is used to improve the speed of spatial computations by constraining search to a portion of the database.

```

george washington university (foggy bottom campus)
Role ID: 'UNIV4'      Role pointer 00
Role name: 'university' subrole: 'university campus'
Role class: 'residential' type: 'aggregate-conceptu'
Role deriv: 'terminal-interaction'
Role mark: 'none'
3D Role ID: 'D3ID150' 3D Role pointer 00
27 Points Generic name: 'dc38617' Feature type: 'area'
max coord: 1at N38 54 8 (521) lon W77 3 8 (126)
min coord: 1at N38 53 50 (16) lon W77 2 42 (572)
The length is 308.24 meters.
The width is 271.37 meters.
The display type is 'fill-polygon'.
The area is 112801.81 square meters.
The compactness is 0.03.

george washington university (memorial gym)
Role ID: 'BULD61'      Role pointer 00
Role name: 'building' subrole: 'gymnasium'
Role class: 'recreational' type: 'physical'
Role deriv: 'hand-segmentation'
Role mark: 'none'
3D Role ID: 'D3ID172' 3D Role pointer 00
composition: 'concrete'
status: 'active'
19 Points Generic name: 'dc38617' Feature type: 'area'
max coord: 1at N38 53 53 (211) lon W77 3 0 (238)
min coord: 1at N38 53 50 (722) lon W77 2 57 (572)
The length is 66.23 meters.
The width is 55.60 meters.
The height is 16 meters.
The display type is '3D-structure'.
The area is 2863.40 square meters.
The compactness is 0.06.

george washington university (gwu medical center)
Role ID: 'BULD129'      Role pointer 00
Role name: 'building' subrole: 'hospital'
Role class: 'residential' type: 'physical'
Role deriv: 'hand-segmentation'
Role mark: 'none'
3D Role ID: 'D3ID310' 3D Role pointer 00
composition: 'concrete'
status: 'active'
15 Points Generic name: 'dc38618' Feature type: 'area'
max coord: 1at N38 54 1 (729) lon W77 3 6 (750)
min coord: 1at N38 53 58 (662) lon W77 3 2 (369)
The length is 72.63 meters.
The width is 67.85 meters.
The height is 18 meters.
The display type is '3D-structure'.
The area is 4468.87 square meters.
The compactness is 0.03.

george washington university (lerner hall)
Role ID: 'BULD137'      Role pointer 00
Role name: 'building' subrole: 'student union'
Role class: 'educational' type: 'physical'
Role deriv: 'hand-segmentation'
Role mark: 'none'
3D Role ID: 'D3ID319' 3D Role pointer 00
composition: 'concrete' & 'brick'
19 Points Generic name: 'dc38618' Feature type: 'area'
max coord: 1at N38 54 1 (99) lon W77 2 53 (44)
min coord: 1at N38 53 58 (602) lon W77 2 49 (478)
The length is 64.14 meters.
The width is 64.90 meters.
The height is 23 meters.
The display type is '3D-structure'.
The area is 2695.99 square meters.
The compactness is 0.04.

```

Figure 3-6: Multiple Roles for 'george washington university'

In MAPS we perform geometric computations on the feature data in the geodetic coordinate system using point, line, and polygon as map primitives. We constrain search by using a hierarchical representation computed directly from the underlying map data. These spatial constraints can be viewed as natural, that is, intrinsic to the data, and may have some analogy to how humans organize a "map in the head" to avoid search. For example, when a tourist who is looking on a map for the location of the Washington Monument is told that the building is in Northwest Washington, they will not spend much time looking at the portion of the map containing Virginia. Depending on their familiarity with the area, they may avoid looking at much of the map outside of Northwest Washington. If they are also told that the monument is also near the tidal basin, that should further constrain their search. As we begin to represent large numbers of map features with more complex interrelationships, we believe that the use of natural hierarchies in urban areas, such as political boundaries, neighborhoods, commercial and industrial areas, serve to constrain search. They may also allow us to build systems that organize data using spatial relationships that are close to human spatial models.

kennedy center (.)
 Role ID: 'BUL08' Role pointer 00
 Role name: 'building' subrole: 'concert hall'
 Role class: 'cultural feature' type: 'physical'
 Role deriv: 'landmark-description'
 Role mark: 'none'
 3D Role ID: 'D3ID21' 3D Role pointer 00
 1 Point Generic name: 'unknown' Feature type: 'point'
 max coord: lat N38 53 44 (880) lon W77 3 21 (100)
 min coord: lat N38 53 44 (880) lon W77 3 21 (100)
 latitude 38 53 44 880
 longitude 77 3 21 100
 elevation 0 meters
 908.237 in /vise/washdc/wgl/dc38617/1bw.img
 landmark image at resolution 1

kennedy center (concert hall)
 Role ID: 'BUL01' Role pointer 00
 Role name: 'building' subrole: 'performing arts comp'
 Role class: 'cultural feature' type: 'aggregate-phy'
 Role deriv: 'terminal-interaction'
 Role mark: 'none'
 3D Role ID: 'D3ID1' 3D Role pointer 00
 composition: 'concrete' && 'stone'
 status: 'active'
 5 Points Generic name: 'dc38617' Feature type: 'areal'
 max coord: lat N38 53 46 (880) lon W77 3 24 (263)
 min coord: lat N38 53 41 (226) lon W77 3 19 (646)
 The length is 193.01 meters.
 The width is 113.81 meters.
 The height is 24 meters.
 The age is 22 years.
 The display type is '3D-structure'.
 The area is 17127.66 square meters.
 The compactness is 0.06.

kennedy center
 Definition: A performing arts complex in northwest DC. Located on the east bank of the Potomac, adjacent to Theodore Roosevelt Island.

Description: Located at the corner of New Hampshire Ave. and F St., N.W. Dedicated to the performing arts, the center houses a concert hall, opera house, stage theater and film theater. Performances are throughout the year. The National Symphony and American Ballet Theater are the resident symphony and ballet, respectively. Beautiful appointments throughout the building are gifts from all over the world, given in memory of John Kennedy for whom the center serves as a living memorial. The rooftop affords a city-wide view, while the terrace provides a panorama of Virginia just across the Potomac.

The landmark image is oriented with north at the top.
 The height is 30 meters.
 The display type is 'point'.

Figure 3-7: CONCEPT for 'kennedy center'

3.2.1. Hierarchical Containment Tree

The hierarchical containment tree is a tree structure where nodes represent map features. Each node has as its descendants those features that it completely contains in *<latitude/longitude/elevation>* space. The hierarchical tree is initially generated by obtaining an unordered list of features (containment list) for each map database feature. Starting with a designated root node ('greater washington d.c.') which contains all features in the database, descendant nodes are recursively removed from the parent node list if they are already contained in another descendant node. The result is that the parent node is left with a list of descendant features that are not contained by any other node. These descendant nodes form the next level of an N-ary tree ordered by the 'contains' relationship. This procedure is performed recursively for every map feature. Terminal nodes are point and line features, or areal features that contain no other map feature. We will discuss the point containment and closest point computation using the hierarchy tree in the following section.

Figure 3-8 shows a small section of the hierarchical containment tree. The use of conceptual features-- features with no physical realization in the world but which represent well understood spatial areas-- can be used to partition the database. In this

case the map feature 'foggy bottom' allows us to partition some of the buildings and roads that are contained within 'northwest washington'. As more neighborhood areas and city districts are added to our database, we expect to see improved performance especially in areas with dense feature distributions. This will also improve the richness of the spatial description available to the user.

```

41 entries for 'contains' for 'northwest washington'
  entry 0: 'mcmillan reservoir (role: 0)'
  entry 1: 'kennedy center (role: 0)'
  entry 2: 'ellipse (role: 0)'
  entry 3: 'executive office building (role: 0)'
  entry 4: 'white house (role: 0)'
  entry 5: 'treasury building (role: 0)'
  entry 6: 'department of commerce (role: 0)'
  entry 7: 'museum of history and technology (role: 0)'
  entry 8: 'key bridge (role: 0)'
  entry 9: 'thomas circle (role: 0)'
  entry 10: 'dupont circle (role: 0)'
  entry 11: 'foggy bottom (role: 0)'
  entry 12: 'whitehurst freeway (role: 0)'
  entry 13: 'mclean gardens (role: 0)'
  entry 14: 'macomb playground (role: 0)'
  entry 15: 'theodore roosevelt island (role: 0)'
  entry 16: 'interior department (role: 0)'
  entry 17: 'district building (role: 0)'
  entry 18: 'lafayette park (role: 0)'
  entry 19: 'constitution hall (role: 0)'
  entry 20: 'national press building (role: 0)'
  entry 21: '23rd street (role: 0)'
  entry 22: 'constitution avenue (role: 0)'
  entry 23: 'virginia avenue (role: 0)'
  entry 24: 'national zoo (role: 0)'
  entry 25: 'georgetown (role: 0)'
  entry 26: 'glover park (role: 0)'
  entry 27: 'national cathedral (role: 0)'
  entry 28: '21st street (role: 0)'
  entry 29: 'north 20th street (role: 0)'
  entry 30: '19th street (role: 0)'
  entry 31: 'east pennsylvania avenue (role: 0)'
  entry 32: 'e street (role: 0)'
  entry 33: 'treasury place (role: 0)'
  entry 34: 'state place (role: 0)'
  entry 35: '26th street (role: 0)'
  entry 36: 'west pennsylvania avenue (role: 0)'
  entry 37: '16th street (role: 0)'
  entry 38: 'l street (role: 0)'
  entry 39: 'vermont avenue (role: 0)'
  entry 40: '13th street (role: 0)'

11 entries for 'contains' for 'foggy bottom'
  entry 0: 'kennedy center (role: 1)'
  entry 1: 'washington circle (role: 0)'
  entry 2: 'state department (role: 0)'
  entry 3: 'american pharmaceutical association (role: 0)'
  entry 4: 'national academy of sciences (role: 0)'
  entry 5: 'federal reserve board (role: 0)'
  entry 6: 'national science foundation (role: 0)'
  entry 7: 'civil service commission (role: 0)'
  entry 8: 'c street (role: 0)'
  entry 9: '22nd street (role: 0)'
  entry 10: 'south new hampshire avenue (role: 0)'

```

Figure 3-8: MAPS: Hierarchical Spatial Containment

3.2.2. Hierarchical Search

In this section we discuss the use of our hierarchical organisation to partition the map database to improve performance by decreasing search when computing the spatial relationships of map features. The hierarchical searching algorithm is basically an N-ary tree searching algorithm. Consider a user at the CONCEPTMAP image display who points at an island in the image and invokes the geometric database to compute its symbolic description. First, using image-to-map correspondence, the system calculates the following map coordinates:

latitude N 38 53 49 (276)
longitude W 77 03 53 (337)

This point is converted into a temporary map database feature and is tested against the root node of the hierarchy tree. If it is not contained in this node (which is generally not the case), then the point cannot correspond to a database feature, and the search terminates. The user is informed that the point is outside the map database. This can occur by pointing at an image that is outside the coverage area of the database or by user entering arbitrary coordinates through the terminal keyboard. Therefore the database has some crude idea of its extent of map knowledge. If the 'contains' test succeeds, a recursive search of the hierarchy tree is performed, by testing containment against the siblings of the node just tested. The search allows several paths to exist for any point, thus more than one sibling may contain a path to the point. This sort of anomaly occurs when a feature happens to exist in the intersecting region of two larger regions. However, if the feature is not contained by the node, then it cannot be contained by any of the node's descendants, and that portion of the tree is not searched further. Figure 3-9 shows the answer to our hypothetical query. The query point is contained within 'theodore roosevelt island', and two search paths in the containment tree are given. In this example, 'theodore roosevelt island' is wholly contained within the 'potomac river' and 'northwest washington'. Since the 'potomac river' is only partially contained within 'northwest washington', and vice versa, two independent paths through the tree are found. The same mechanism used for point queries is used for line and polygon features, although the primitive determination of containment is specialized for the geometric type of the feature. The ability to perform hierarchical search and to represent *multiple paths* through a spatial database addresses some of the major flaws with pure hierarchical decompositions. In CONCEPTMAP it is the basis for 'level of detail' queries using depth limiting search, described in Section 2.2, by associating a scale attribute with each node in the tree. When the scale of the node is finer than the search scale, search is terminated at that node. The successful search path to that point is returned. We have recently implemented this technique and have used the scale of the imagery being displayed to set the search scale attribute by making some simple assumptions about the ground sample distance of the imagery. It is very interesting to point to an area in a 1:128000 scale image window, perform a hierarchical query, and then perform the same operation in a 1:12000 scale image window. One can observe that the search is significantly faster using the low

```

This node belongs in the following place(s):
database 'washdc' last compiled '9/4/86'
3 entries for 'contains' for 'theodore roosevelt island (Role: 0)'
entry 0:      'northwest washington (Role: 1)'
entry 1:      'district of columbia (Role: 0)'
entry 2:      'greater washington d.c. (Role: 0)'
               ***** A N D *****
2 entries for 'contains' for 'theodore roosevelt island (Role: 0)'
entry 0:      'potomac river (Role: 0)'
entry 1:      'greater washington d.c. (Role: 0)'

```

**Figure 3-9: MAPS: Containment Tree Entry for
Theodore Roosevelt Island**

resolution image (1:128000) and that features reported by the high resolution search are omitted under low resolution. A rigorous analysis of speedup is difficult since it depends primarily on the distribution of features, their spatial relationships, and the location of the query with respect to feature density. However, we observe factors of 30 to 100 speedup over brute-force containment queries using our hierarchical search. An additional speedup factor of 5 to 10 can be achieved using depth limiting search. Using the CONCEPTMAP Washington D.C. database with approximately 250 concepts and 600 roles such queries take between 1 and 5 cpu seconds on a VAX 11/780 with floating point accelerator. In Section 3.4 we describe the use of simultaneous hierarchical queries in multiple databases using this same methodology.

3.3. Unifying Spatial and Factual Queries

In Section 3.1 we described the three levels of attribution available for *concept* and *role* schemata. In this section we describe the user's view of these attributes via database query. As depicted in Figure 3-10, there are three broad classes of queries available to users within CONCEPTMAP. *Spatial queries* can be performed by specifying a *concept* or *role* schema, or by indicating an area of interest on an image display. *Factual queries* are specified by constructing a *query schema* by specifying one or more slot values for any of the three classes of *role* slots. Some representative attributes are given in Figure 3-11. Both exact and partial matching can be performed between the constructed query schema and CONCEPTMAP role schemata.

The result of *spatial* and *factual* queries is the production of a *query-list* schema. This *query-list* contains a list of *role-id* values that satisfied the query. In the case of spatial queries, it may contain geometric facts related to the query, such as the actual computed points of intersection, adjacency, or minimal distance. The third class of queries is logical operations on these *query-list* schemata. They include *logical-and*, *logical-or*, *logical-not* and *logical-exclusive-or*. In addition, operations to query on role schema that are

Spatial	Factual	Query-list Set
-----	-----	-----
subsumed by	exact template match	logical and
intersection	partial template match	logical or
closest		logical exclusive or
adjacent		logical not
in-radius		
contains		

Figure 3-10: Query Operations on Role Schemata

principal role for concepts, to merge *query-list* schemata, and to enumerate all roles given the concept *principal role* are available. The user's view of spatial attributes is uniform and hides the actual implementation of each attribute class; all attributes in Figure 3-11 are available for *factual* queries. Combinations of *spatial* and *factual* queries are possible by performing each query and performing the appropriate *query-list* set algebra. The CONCEPTMAP query interface can name and save *query-list* schemata for future use. An example of this is *memo-files* which cache previously computed spatial relationships performed over the entire database. The user can choose to use the *memo-files* or compute these relationships dynamically. However, for the most part, application programs rather than interactive human users make use of the *memo-file*. A case in point is the SPAM system (section 4.4) which stores spatial relationships developed during image analysis.

System Defined	User Defined	
-----	-----	
rolename	composition	
roleclass	status	
roletype	class	
rolederivation	activity	
subrolename		
Property List Descriptions		
-----	-----	-----
length	width	height
radius	age	capacity
stories	basement	population
display-type	multi-family	single-family
area	compactness	fractional fill
scale	perimeter	color
shape		

Figure 3-11: Template Match Query Attributes

Collections of CONCEPTMAP databases can be searched to satisfy spatial queries. The hierarchy tree associated with each CONCEPTMAP database contains the spatial extent of the database and can be used to quickly detect whether a query can potentially be satisfied by searching the tree. Entire trees can therefore be pruned from consideration using a simple test for spatial overlap. Since the *concept-id* and *role-id* are uniquely assigned by the system, concepts and roles are not restricted to a particular database, they can be replicated or referenced by making a copy of the identifier and can be uniquely determined by referencing the identifier within the context of a named CONCEPTMAP database. Figure 3-12 shows how concepts can be shared among databases. The databases USA and WASHDC contain the concepts 'virginia', 'maryland', and 'district of columbia'. The concept 'national airport' is represented in WASHDC and NATAIR. Within the WASHDC database 'national airport' is represented by its principal *role* while the NATAIR database contains *role* schema for runways, terminal buildings, hangars, and other features at the airport.



- Large monolithic databases can be partitioned to improve average retrieval performance.
- Databases can be physically disjoint, i.e., located on several workstations connected by a local area network.
- Databases can be specialized for particular tasks by grouping *concepts* and *roles* with common properties.

One obvious problem is that of data consistency. In order to insure consistency of *concepts* across multiple databases we limit database modification to the addition, modification, and deletion of non-principal *roles*. That is, the *concept-name*, *concept-id* and *principal role* cannot be changed unless all instances of the *concept* are deleted. However, the number of non-principal *roles* associated with the *concept* can vary, as shown in the previous example. Some of these restrictions can be relaxed by employing a more sophisticated transaction monitor and locking mechanism.

An example of the utility of spatial queries across multiple databases is illustrated by the following example. Consider a spatial query to the CONCEPTMAP database using the geographic location of the main terminal building at National Airport in Washington D.C..

```
latitude N 38 50 55 (628)
longitude W 77 02 31 (347)
```

Using multiple databases, USA, WASHDC, and NATAIR the containment trees in Figure 3-13 are generated. In the context of the USA database the query point is found as within the concept 'alexandria', within 'national airport (Role 1)' in the context of WASHDC, and within 'national airport (Role 10)' in database NATAIR. Depending on the specific task, this basic information can be used to query and infer relationships across databases such as: Is 'national airport' contained within 'alexandria'? Additionally, multiple database search allows applications to locate a highly detailed database with limited area coverage (NATAIR) within a larger context (WASHDC and USA) without specifically generating those relationships or generating and searching a single large monolithic database.

```
This node belongs in the following place(s):
database 'usa' last compiled '8/24/86'
2 entries for 'contains' for 'alexandria (Role: 3)'
entry 0:      'virginia (Role: 0)'
entry 1:      'usa (Role: 0)'

database 'washdc' last compiled '9/4/88'
3 entries for 'contains' for 'national airport (Role: 1)'
entry 0:      'national airport (Role: 0)'
entry 1:      'virginia (Role: 0)'
entry 2:      'greater washington d.c. (Role: 0)'

database 'natair' last compiled '4/11/86'
3 entries for 'contains' for 'national airport (Role: 10)'
entry 0:      'national airport (Role: 9)'
entry 1:      'national airport (Role: 1)'
entry 2:      'national airport (Role: 0)'
```

Figure 3-13: Search In Multiple Databases

4. Application of MAPS In Computer Vision and Graphics

In the previous section we have discussed the structure of the CONCEPTMAP database and some methods to query and access spatial data. In this section we focus on the use of MAPS by application programs primarily designed to perform analysis and interpretation of aerial imagery. The interpretation of remotely sensed data requires knowledge about the scene under consideration. Knowledge about the type of scene (airport, suburban housing development, urban city) can aid in low-level and intermediate-level image analysis, and can be expected to drive high-level interpretation by constraining search for plausible consistent scene models. Loosely speaking, this knowledge base should contain known facts and spatial relationships between objects in an area of interest, access to historical or *a priori* map knowledge, and methods to relate earth coordinates to pixel locations in digital imagery. Unfortunately, these spatial database capabilities are somewhat different than those found in traditional geographic information systems. Other issues include methods for spatial knowledge utilization and representation. For example, simply having access to cartographic descriptions does not really address the problem of how to operationalize iconic descriptions for image analysis and interpretation. One can identify several explicit requirements for such a spatial database system.

- Model spatial knowledge explicitly at multiple levels of detail.
- Provide access to factual queries using partial matching.
- Provide access to imagery and spatial entities using spatial constraints.

4.1. A Geodetic Frame Of Reference

An implicit requirement crucial to successful operationalization of spatial knowledge is that the metrics used by the analysis system be defined in cartographic coordinates, such as *<latitude/longitude/elevation>*, rather than in an image-based coordinate system. Systems that rely on descriptions such as "the runway has area 12000 pixels" or "oil storage tanks are between 212 and 345 pixels in area" are useless except for (perhaps) the analysis of one image. Further, spatial analysis based on the semantics of *above*, *below*, *left-of*, *right-of* etc., are also inappropriate for general interpretation systems. To operationalize metric knowledge one must relate the world model to the image under analysis. This should be done through image-to-map correspondence using camera models. We can directly measure ground distances, areas, absolute compass direction, and recover crude estimates of height using a camera model computed for each image under analysis. The MAPS system allows us to express spatial knowledge in ground based metrics, which in turn allows us to work with imagery and map data at a variety of scales and resolutions.

4.2. Rule-Based Systems In Image Analysis

The primary role for rule-based systems in the analysis of remotely sensed imagery is to provide constraints so that image analysis and interpretation tools can be used in spite of their inherently errorful performance. The goal then, is to integrate rule-based systems with image analysis techniques to constrain the search space of possible scene interpretations using domain and general knowledge. These constraints can be characterized as *"what to look for"* and *"where to look for it"*. Currently, rule-based systems are most powerful (and successful) in narrow, well-defined task areas. Our research has been focused on tasks such as road finding and tracking, urban scene analysis, and airport scene analysis. These are narrow segments of the much larger problem. However, as most researchers in this area recognize, we can only achieve sparse performance data points in a large task space by building, measuring, and evaluating performance systems. It is certainly the case that in order to solve the general remote sensing problem an analysis system will require utilization of general problem solving capabilities and vast amounts of domain and common sense knowledge, currently far beyond the capability of any research system.

However, even without general problem solving capabilities, there is much to be gained by the development of "existence proof" rule-based interpretation systems. Current image processing systems are incapable of high-level description of the results of image interpretation. The combination of a map-based world model and rule-based systems which have site-specific or task-specific knowledge can be used to bridge the gap between users and current state-of-the-art image interpretation systems. The long-term goal of our research is to develop systems which can interact with a human analyst at a highly symbolic level, which maintain a world database of previous events, and use expert level knowledge to predict areas for fruitful analysis, and integrate the results of the analysis into a coherent model. As we experiment with such systems it is important to explore the effects and utility of various knowledge sources on the quality of the final image analysis. Too often, this empirical approach is overlooked, and it is difficult to evaluate 'where the action is' in any particular system. This also implies that researchers utilize system architectures where particular sets of knowledge or extraction techniques can be easily added or removed for experimentation and measurement.

In the following sections we will briefly describe several image analysis and interpretation systems that utilize components of the MAPS system for the representation of world knowledge, spatial constraints, for the selection of imagery based upon geodetic coverage and scale, and for computer image generation.

4.3. Map-Guided Image Segmentation

MACHINESEG¹⁰ is a program that performs map-guided image segmentation. It uses map knowledge to control and guide the extraction of man-made and natural features from aerial imagery using a region-growing image analysis technique. MACHINESEG uses the CONCEPTMAP database from the MAPS system as its source of map knowledge. Figure 4-1 shows the interaction between the map database and image processing and feature extraction tools. Map knowledge can be used to represent generic shapes, sizes, and spectral properties typical to a large class of objects such as roads, or can describe specific features such as known buildings where geodetic position may be known as well as the more general structural properties. In the latter case we have been able to segment a diverse set of cultural features such as buildings, reservoirs, and roads using map-to-image correspondence to project map-based descriptions onto a new image under analysis. This projection generally only provides a coarse idea of the actual position of the feature, but greatly constrains search in the image. MACHINESEG uses the following components of the MAPS system.

- CONCEPTMAP to retrieve shape and position models.
- Map-to-image correspondence to project models onto new imagery and image-to-map correspondence to calculate metric distances and areas.
- CONCEPTMAP to store extracted features from several images prior to interactive editing and integration into the database.

The notion of map-guided image segmentation is not a new one. Many researchers have discussed the use of *a priori* knowledge of various object features such as size, shape, orientation, and color to extract and identify features from an image. However, there are few, if any, examples of systems that can systematically search through a database of images looking for examples of particular objects or classes of objects.

It is important to characterize what we mean by "map-guided" image segmentation. Map-guided image segmentation is the application of task-independent spatial knowledge to the analysis of a particular image using an explicit map-to-image correspondence derived from camera and terrain models. Map-guided segmentation is not interactive editing or computation of descriptions in the image domain, since these descriptions are valid only for one specific image. As we have described, each role schema contains a *role-geographic-id* which points to a geodetic description (*<latitude,longitude,elevation>*) for each map entity in the CONCEPTMAP database. This description is in terms of *points*, *lines*, and *polygons*, or collections of these primitives. Features such as buildings, bridges, and roads have additional attributes describing their elevation above the local terrain, as well as their composition and appearance. The location of each map feature in the database can be projected onto a new image using a map-to-image correspondence maintained by MAPS. Likewise, a new map

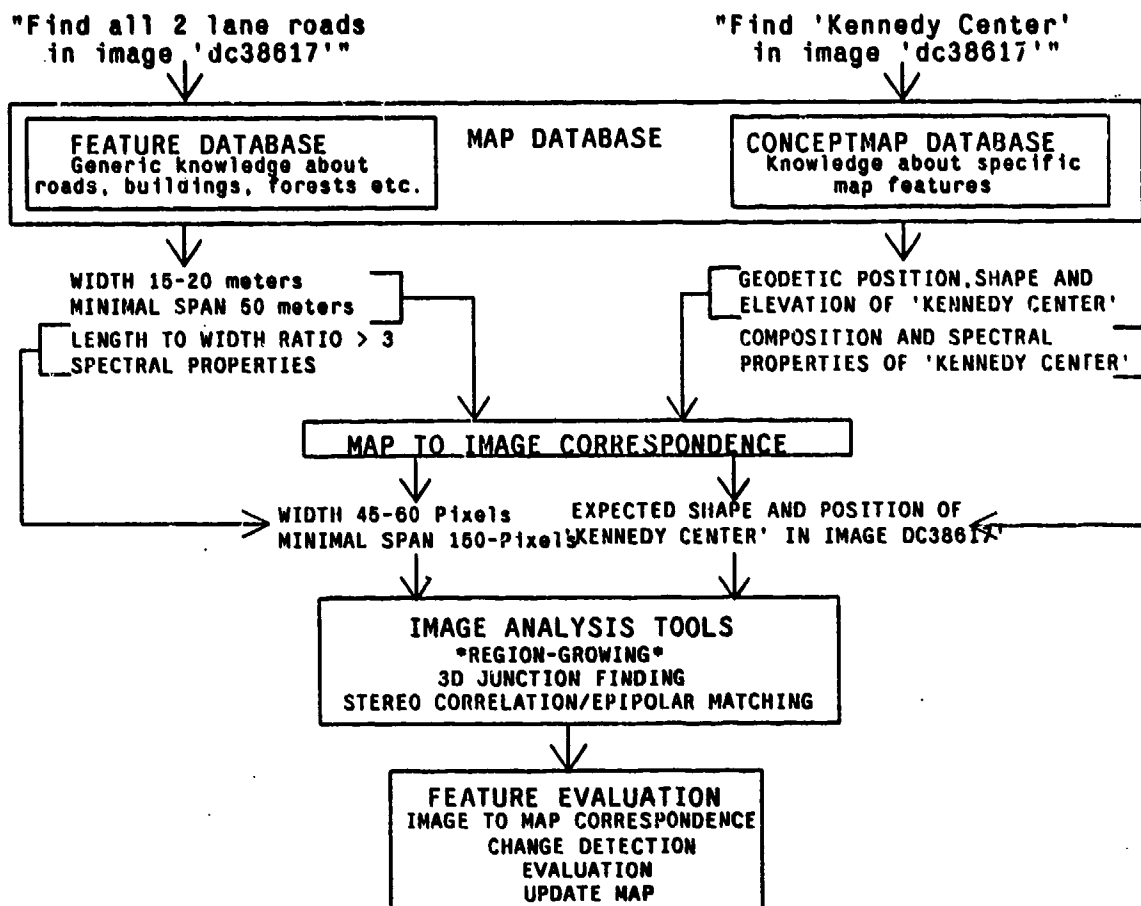


Figure 4-1: MACHINESEG: Map-Guided Feature Extraction

feature can be projected onto the existing image database. If camera model errors are known, one can directly calculate an uncertainty for image search windows. Further, as new features are acquired their positions can be directly integrated into the map database using image-to-map correspondence procedures.

Figure 4-1 shows a schematic description of the map-guided feature extraction process in MAPS. There are two methods for applying map knowledge to the extraction of features from aerial imagery. The first method uses *generic knowledge* about the shape, composition and spectral properties of man-made and natural features. This may be provided by the knowledge base of the application. The second uses *map-based template* descriptions. These descriptions are stored in the CONCEPTMAP database and represent knowledge about known buildings, roads, bridges, etc. This knowledge includes geodetic position, shape, elevation, composition and spectral properties. In the second case, the position, orientation, and scale are constrained whereas in the first, only the scale can be determined. In both cases, in order to operationalize spatial knowledge for the analysis of

a particular image, a map-to-image correspondence is performed.

4.4. Airport Scene Interpretation

SPAM, System for Photo interpretation of Airports using MAPS, is an image-interpretation system. SPAM^{11, 12} coordinates and controls image segmentation, segmentation analysis, and the construction of a scene model. It provides several unique capabilities to bring map knowledge and collateral information to bear during all phases of the interpretation. These capabilities include:

- The use of domain-dependent spatial constraints to restrict and refine hypothesis formation during analysis.
- The use of explicit camera models that allow for the projection of map information onto the image.
- The use of image-independent metric models for shape, size, distance, absolute and relative position computation.
- The use of multiple image cues to verify ambiguous segmentations. Stereo pairs or overlapping image sequences can be used to extract information or to detect missing components of the model.

Figure 4-2 shows the overall organization of the interpretation system. SPAM maintains an internal spatial database that is composed of feature extracted from imagery by various methods, possibly from several images, where the features are represented in terms of their geodetic position rather than their image coordinates. In fact, SPAM performs interpretation in map-space which allows for a variety of knowledge such as maps and multi-temporal imagery to be handled in a uniform manner.

SPAM uses the following components of the MAPS system.

- CONCEPTMAP to retrieve shape and position models and site-specific map knowledge.
- Map-to-image correspondence to project models onto new imagery and image-to-map correspondence to calculate metric distances and areas.
- Procedures to compute spatial relationships between hypotheses including containment, intersection, adjacency, closest point of approach, subsumed by.

4.5. Stereo Image Analysis

STEREOSYS¹³ is a flexible stereo verification system. Stereo verification refers to the verification of hypotheses about a scene by stereo analysis of the scene. Unlike stereo interpretation, stereo verification requires only coarse indications of three-dimensional structure. In the case of aerial photography, this means coarse indications of the heights of objects above their surroundings. STEREOSYS is used within the SPAM system to

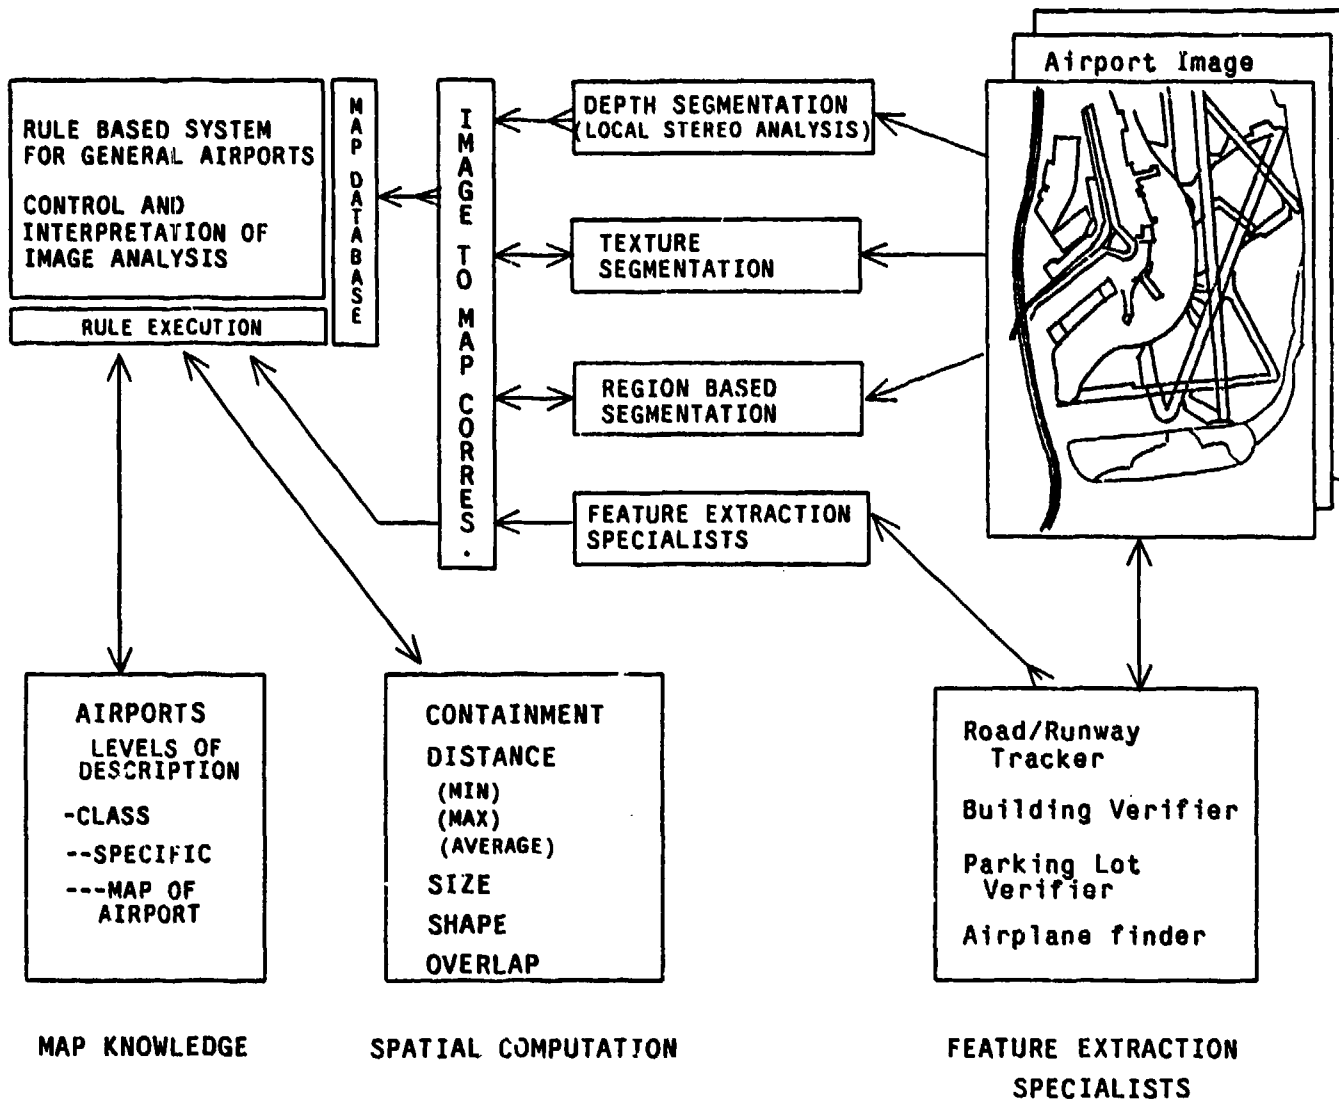


Figure 4-2: SPAM: System Organization

confirm or refute airport feature hypotheses based upon their three-dimensional structure.

Stereo verification deals with a variety of problems that are not ordinarily present in isolated experiments with stereo matching and analysis. Some of the most interesting problems within the spatial database context are the following:

- How to select an appropriate conjugate image pair from a database of overlapping images based on criteria that would maximize the likelihood for good correspondence.
- The image pairs must be dynamically resampled such that the epipolar assumption (i.e., epipolars are scan lines) used in most region-based stereo

matching algorithms can be applied.

- An initial coarse registration step is generally necessary because the quality of the correspondence between conjugate pairs varies greatly. In many cases the magnitude of the initial misregistration is greater than the expected disparity shift.

These requirements, in turn, raise a broad set of research issues. In terms of spatial databases the major questions are related to how an aerial image database can be used to automatically generate a useful stereo pair containing an arbitrary region and how a stereo system can handle the misregistration problems inherent in multi-source image databases. The results of this research indicate that image/map database issues in stereo verification influence the utility of such an approach as much as the underlying stereo matching algorithm. In fact, they are intimately related. The ability to be flexible in the selection of stereo pairs provides opportunities for multi-temporal, multi-scale, or multi-look matching. Equally important is flexibility in the matching algorithm, especially with respect to assumptions that require nearly perfectly aligned conjugate images.

Figure 4-3 shows the overall organization of the verification system. We believe that the ability to dynamically select conjugate image pairs from a database based upon the region of interest and knowledge of the requirements of the matching algorithm is required for a fully automated image analysis system.

STEREOSYS uses the following components of the MAPS system.

- IMAGE database to select appropriate conjugate pair imagery based on time, scale and flightline information.
- Image-to-map correspondence to resample imagery to a common rectified projection so that epipolars align as scan lines in the image.
- TERRAIN database for resampling imagery.

4.6. Computer Image Generation

Computer graphics plays an increasingly important role in the general areas of image processing, photo-interpretation and cartography. We have constructed an interactive graphics system, WASH3D¹⁴ that integrates information extracted from aerial imagery, digital terrain databases, and cultural feature database and allows a user to generate cartographically accurate color 3D perspective scene displays of the Washington, D.C., area. Figure 4-4 depicts the steps in the generation of a 3D scene. Opportunities exist for many types of spatial data fusion and generation of specialized scenes simply by the appropriate selection of subsets of features stored in the MAPS database. The advantages of a unified database as viewed by the user or application program appear to be very high when compared to the alternative of specialization within the database for each display

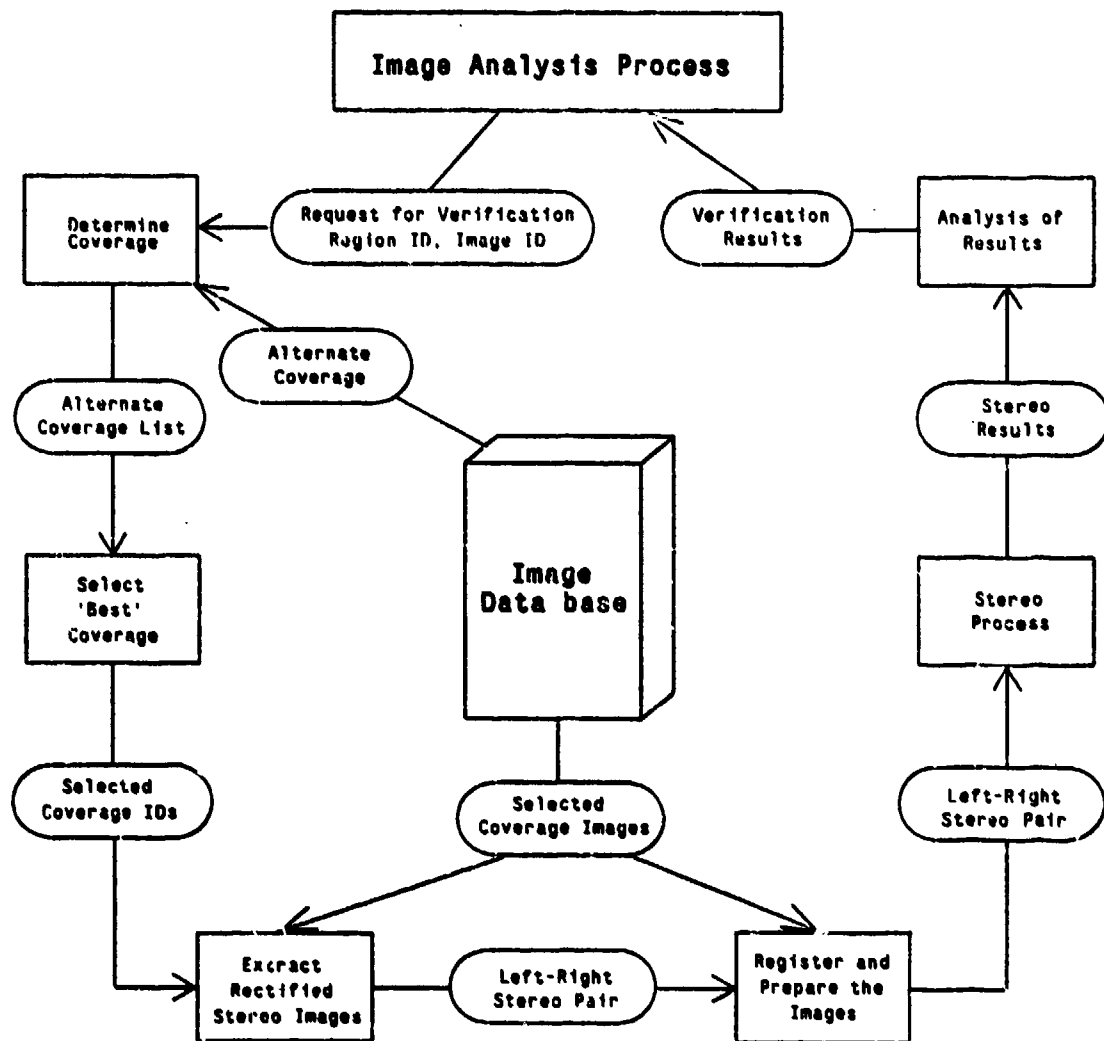


Figure 4-3: STEREOSYS: Stereo Verification

task. WASH3D uses the following components of the MAPS system.

- TERRAIN database for elevation data and interpolation to arbitrary scene resolution.
- DMA DLMS database for large scale cultural features and terrain coverage providing coarse surface material thematic map.
- CONCEPTMAP database for detailed three-dimensional features such as

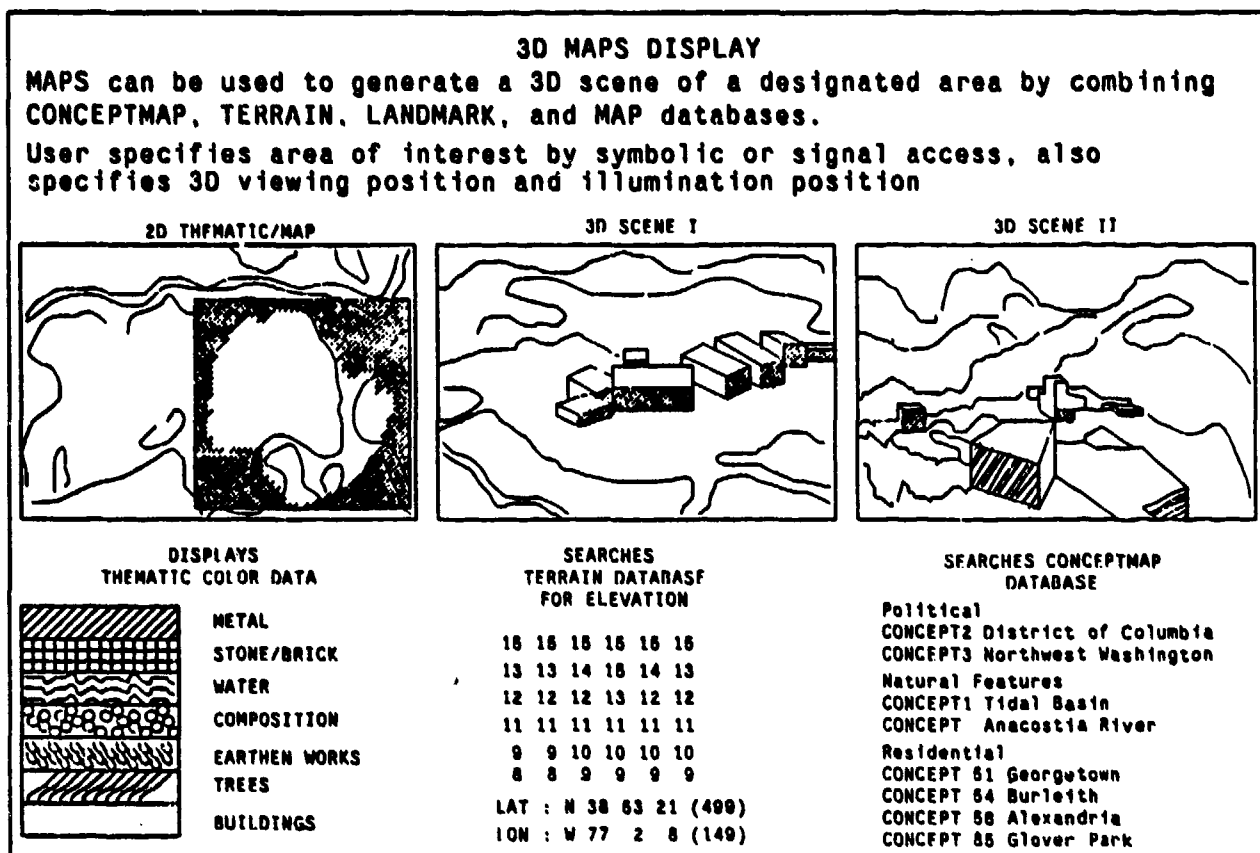


Figure 4-4: WASH3D: Computer Image Generation

buildings, roads, bridges, and small-scale cultural features omitted from the DLMS database.

- LANDMARK database for drawing significant point features.

Photograph 4-5 shows an overview of the Washington D.C. area viewed from the south with Theodore Roosevelt Island in the center of the scene. The surface coloring is a combination of the surface material polygons in the DLMS database and transparency overlay of concepts from the CONCEPTMAP database. For example, the Foggy Bottom area, Reflecting Pool, and White House areas are generated from the CONCEPTMAP database, as they do not appear in the baseline DLMS data. Photograph 4-6 shows the same area as Photograph 4-5 except that three-dimensional features such as buildings, roads, and bridges from the CONCEPTMAP database have been rendered, and the transparency overlay has been omitted, giving a coarser impression of surface terrain. Photograph 4-7 shows a closeup of the Foggy Bottom area with buildings and roads rendered and with the same enhanced terrain characterization as in Photograph 4-5.

Figure 4-5: WASH3D: Terrain Scene of Washington D.C. With CONCEPTMAP Database Overlay




Figure 4-6: WASII3D: Scene of Washington D.C. With CONCEPTMAP
Cultural Features



Figure 4-7: WASH3D: Closeup of Foggy Bottom Area With CONCEPTMAP
Terrain and Cultural Features



5. Conclusions

In this paper we described some thoughts on how Artificial Intelligence might play a role in the development of enhanced capabilities for spatial databases. We proposed the extension of the concept of a geographic information system to actively include remotely sensed data along with image-to-map correspondence capabilities necessary to relate images, maps, terrain, and external map data into a common geodetic-based representation. We described a research system, MAPS, which has been developed to explore these ideas. A schema representation for spatial data, methods to constrain search using hierarchical containment trees, depth and breadth limiting search, and the concept of level-of-detail representations are introduced. Several knowledge-based image interpretation systems developed at Carnegie Mellon are used as examples of how spatial knowledge can be utilized. Specifically, those components of the MAPS system that were particularly relevant to their overall system design or task performance were described.

6. Acknowledgements

Dan Edwards and John McDermott commented on an early draft of this paper. George Lukes, Philip Swain, and the referees provided excellent comments on the final draft. Robert Lai aided in the preparation of many of the diagrams. The research described in this paper was performed by members of the Digital Mapping Laboratory, Paul Carnine, Jerry Denlinger, Wilson Harvey, Bruce Irvin, John Kraft, Robert Lai, Dave McKeown, Cliff McVay, and Lambert Wixson.

7. References

1. *Geographic Information Systems in Government*, A. Deepak Publishing, Hampton, VA, 1986.
2. Michael L. Brodie and John Mylopoulos, editor, *On Knowledge Base Management Systems*, Springer-Verlag Inc., Topics in Information Systems, 1986.
3. Buchanan, B.G., "EXPERT SYSTEMS: Working Systems and the Research Literature", Tech. report STAN-CS-86-1094, Stanford University, December 1985.
4. Donald A. Waterman, *A Guide To Expert Systems*, Addison-Wesley Publishing Company, Reading, MA, 1986.
5. John McDermott, "Making Expert Systems Explicit", *Information Processing 86*, H.J. Kugler, ed., Elsevier Science Publishers B.V., 1986, pp. 539-544.
6. McKeown, D.M., "MAPS: The Organization of a Spatial Database System Using Imagery, Terrain, and Map Data", *Proceedings: DARPA Image Understanding Workshop*, June 1983, pp. 105-127, Also available as Technical Report CMU-

CS-83-136

7. McKeown, D.M., "Digital Cartography and Photo Interpretation from a Database Viewpoint", in *New Applications of Databases*, Gargarin, G. and Golembe, E., ed., Academic Press, New York, N. Y., 1984, pp. 19-42.
8. McKeown, D. M., "Knowledge-Based Aerial Photo Interpretation", *Photogrammetria, Journal of the International Society for Photogrammetry and Remote Sensing*, Vol. 39/1984, pp. 91-123, Special Issue on Pattern Recognition
9. Avron Barr and Edward A. Feigenbaum, *The Handbook of Artificial Intelligence*, William Kaufmann, Inc., Los Altos, CA, Vol. 1, 1981.
10. McKeown, D.M., Denlinger, J.L., "Map-Guided Feature Extraction from Aerial Imagery", *Proceedings of Second IEEE Computer Society Workshop on Computer Vision: Representation and Control*, May 1984, Also available as Technical Report CMU-CS-84-117
11. McKeown, D.M., and McDermott, J., "Toward Expert Systems for Photo Interpretation", *IEEE Trends and Applications '89*, May 1983, pp. 33-39.
12. McKeown, D.M., Harvey, W.A. and McDermott, J., "Rule Based Interpretation of Aerial Imagery", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-7, No. 5, September 1985, pp. 570-585.
13. McKeown, D.M., McVay, C.A., and Lucas, B. D., "Stereo Verification In Aerial Image Analysis", *Optical Engineering*, Vol. 25, No. 3, March 1986, pp. 333-346, Also available as Technical Report CMU-CS-85-139
14. McKeown, D. M., and J. L. Denlinger, "Graphical Tools for Interactive Image Interpretation", *Computer Graphics*, Vol. 16, No. 3, July 1982, pp. 189-198.